

UNITED STATES PATENT APPLICATION

**SERVICE SELECTION IN A SHARED ACCESS NETWORK
USING TUNNELING**

INVENTORS:

John W. Garrett

Charles Robert Kalmanek Jr.

Han Q. Nguyen

Kadangode K. Ramakrishnan

Cross Reference to Related Applications

This application claims priority to United States Provisional Application Serial No. 60/190,633, entitled "INTERNET SERVICE SELECTION OVER CABLE," filed on March 20, 2000, and to United States Provisional Application Serial No. 60/190,636, entitled "QUALITY OF SERVICE OVER THE HFC CABLE PLANT," filed on March 20, 2000, the contents of which are incorporated by reference herein.

100 }
101 }
102 }
103 }
104 }
105 }
106 }
107 }
108 }
109 }
110 }
111 }
112 }
113 }
114 }
115 }
116 }
117 }
118 }
119 }
120 }
121 }
122 }
123 }
124 }
125 }
126 }
127 }
128 }
129 }
130 }

SERVICE SELECTION IN A SHARED ACCESS NETWORK USING TUNNELING

Field of the Invention

5 The present invention relates generally to communication network services, and, more particularly, to providing multiple services in a communication network.

Background of the Invention

10 Customers of communication network services often desire access to a plurality of different services and different service providers. For example, when using a dial-up connection to a packet-switched data network such as the Internet, a customer can choose from multiple service providers by dialing different telephone numbers in the PSTN. The physical path from the customer to
15 the customer's Internet Service Provider (ISP) is dedicated to the connection for the duration of the telephone call. The ISP assigns an IP address to the customer and can link the authenticated customer and the assigned IP address to the physical address (e.g. dial-up modem) used by the customer. With this linkage, the ISP can ensure the customer only uses the address authorized by the ISP and
20 can use the customer's IP address to manage access to the ISP's services. The physical connection between a customer and the ISP, as well as the linkage to IP address assignment and customer authentication is terminated when the dial-up connection is terminated.

 Constrained by the physical capacity of these temporary
25 connections across the PSTN, many service providers are moving to high-speed access architectures (e.g., digital subscriber line (DSL), wireless, satellite, or cable) that provide dedicated physical connectivity directly to the subscriber and under the control of the ISP. These alternatives to shared access through the switched telephone network, however, do not lend themselves to shared access by
30 multiple services and/or service providers.

Summary of the Invention

It is an object of the invention to enable multiple services or service providers to share the facilities of an access network infrastructure providing physical connectivity to subscribers. In accordance with an embodiment of the invention, each network access device is assigned two network addresses: a first network address associated with a particular service or service provider to which the user of the device is subscribed and a second network address utilized in the access network infrastructure to tunnel to the relevant service network. Packets from the network access device are encapsulated and routed through the access network infrastructure to arrive at a network node within the associated service network where it is de-encapsulated and routed to its destination. The network access device advantageously may be used in communication network services with a service or service provider that is separate from the operator of the access network infrastructure.

These and other advantages of the invention will be apparent to those of ordinary skill in the art by reference to the following detailed description and the accompanying drawings.

Brief Description of the Drawings

FIG. 1 illustrates an interconnection of packet-switched service networks and an access network embodying principles of the invention.

FIG. 2A and FIG. 2B is conceptual representation of an example embodiment using layer three tunneling illustrating principles of the invention based on an HFC access architecture with corresponding end-to-end protocol layers.

FIG. 3A and FIG. 3B is conceptual representation of another example embodiment using layer two tunneling illustrating principles of the invention based on an HFC access architecture with corresponding end-to-end protocol layers.

FIG. 4 is a conceptual representation of IP encapsulation within IP.

FIG. 5 is a conceptual representation of minimal encapsulation within IP.

FIG. 4 is a flowchart of processing performed at a network access device, in accordance with an embodiment of the invention.

FIG. 5 is a flowchart of processing performed at a tunneling router in the service network, in accordance with an embodiment of the invention.

Detailed Description

In FIG. 1, a plurality of subscribers operating network access devices 101, 102, 103, ... 104 are provided access to communication network services, which are facilitated by a plurality of packet-switched data networks, shown in FIG. 1 as 151 and 152. Packet-switched data networks 151 and 152, referred to herein as "service networks," offer access to different services and/or are operated by different service providers. For example, service network 151 could provide packet-switched connectivity to public data networks while service network 152 could offer packet-switched telephony service (or the same public data network connectivity, but from a different service provider). The service networks, as is well known in the art, utilize a network addressing scheme to route datagrams to and from hosts: for example, where the service networks utilize the TCP/IP protocol suite, Internet Protocol (IP) addresses are assigned to each host and utilized in the process of routing packets from a source to a destination in the networks. See, e.g., "INTERNET PROTOCOL," IETF Network Working Group, RFC 791 (September 1981); S. Deering, R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification," IETF Network Working Group, RFC 1883 (December 1995), which are incorporated by reference herein. The invention shall be described herein with particular reference to the TCP/IP protocol suite and IP addresses, although those skilled in the art would readily be able to implement the invention using any of a number of different communication protocols.

The network access devices 101 ... 104 are typically customer premises equipment (CPE) such as a personal computer, information appliance, personal data assistant, data-enabled wireless handset, or any other type of device

capable of accessing information through a packet-switched data network. Each network access device 101 ... 104 is either connected to or integrated with a

95 network interface unit 111 ... 114, e.g. a modem, which enables communication through an access network infrastructure, shown as 120 in FIG. 1. The access network infrastructure 120 advantageously can be operated and maintained by an entity that is the same as or different from the entities operating and maintaining the service networks 151 and 152. The network access devices 101 ... 104, in

100 accordance with an aspect of the invention, communicate with the packet-switched service networks 151 and/or 152 by what is known in the art as "tunneling." Tunneling is the process by which a packet is encapsulated within another packet, which is delivered between two endpoints of the "tunnel" as a means to alter the conventional routing of the packet. The encapsulated packet

105 travels to an intermediate destination that otherwise would not have been selected based on the destination address indicated in the encapsulated packet. Thus, service-related network traffic is tunneled between the network access devices 101 ... 104 used by the service subscribers and the service networks 151, 152 providing the relevant services. Each network access device in FIG. 1 is assigned

110 at least two IP addresses: (1) an IP address allocated from the address space of the access network infrastructure (this address is used when tunneling to the relevant service network) and (2) an IP address allocated from an address space associated with the particular service or service provider to which the user of the device is subscribed. For example, and for purposes of the description herein, network

115 access device 101 is assumed to have been assigned an IP address associated with the service provider operating service network 151 and an IP address for tunneling to service network 151.

As described in further detail herein, tunneling can be accomplished in different layers of the protocol stack. In one embodiment of the

120 present invention, the technique of IP encapsulation can be utilized—so that the different IP-based services offered by the different service networks 151 and 152 utilize shared layer one and layer two resources in the access network infrastructure 120. FIG. 2A shows an exemplary access architecture for practicing

this embodiment based on a hybrid fiber coaxial (HFC) access network. As is
 125 known in the art, each network interface device 201 ... 202 is either connected to
 or integrated with a cable modem 211 which enables communication through the
 HFC network 221. In accordance with the Data Over Cable Service Interface
 Specification (DOCSIS), a Cable Modem Termination System (CMTS), shown as
 225 in FIG. 2A, communicates with the cable modems 211 and manages access to
 130 both upstream and downstream cable capacity on the HFC networks 221. See,
 e.g., "Data-Over-Cable Service Interface Specifications: Cable Modem
 Termination System – Network Side Interface Specification," Cable Television
 Laboratories, Inc., SP-CMTS-NSI-I01-960702; "Data-Over-Cable Service
 Interface Specifications: Cable Modem to Customer Premise Equipment Interface
 135 Specification," Cable Television Laboratories, Inc., SP-CMCI-C02C-991015;
 "Data-Over-Cable Service Interface Specifications: Baseline Privacy Plus
 Interface Specifications," Cable Television Laboratories, Inc., SP-BPI-I06-
 001215, which are incorporated by reference herein. The CMTS 225 manages
 the scheduling of both upstream and downstream transmission and allocates cable
 140 capacity to individual customers identified by a Service ID (SID). The CMTS
 225 can have an integrated router 228 or can be a separate device 226 that bridges
 to a fast Ethernet switch 227 which connects to the router 228. The IP router 228
 provides connectivity to an IP network 222 which interfaces to IP routers 241 and
 242 in service networks 251 and 252, respectively. Accordingly, the HFC
 145 network 221, the CMTS 225, and the IP network 222 correspond to the access
 network infrastructure 120 shown in FIG. 1. FIG. 2B shows a conceptual diagram
 of the end-to-end communication protocol stack from a network access device 201
 (101) to a router 241 (141) in service provider's network 251 (151) where IP
 encapsulation is being utilized. As is known in the art, the lowest layer deals with
 150 the physical layer (PL) of the protocol stack, e.g. the Ethernet physical media
 device (PMD) layer; the second layer deals with the data link layer, e.g. the
 Ethernet Media Access Control (MAC) layer; while the third layer in the protocol
 stack deals with the network layer, e.g. the IP layer. As shown in the network
 layer of the protocol stack in FIG. 2B, IP traffic between the network access

155 device 201 and the router 241 in the service network 251 is encapsulated within another IP layer.

IP encapsulation may be accomplished using a variety of known techniques. FIG. 4 is a conceptual representation of the process of encapsulating a standard IPv4 packet, in accordance with the technique of "IP Encapsulation Within IP." See C. Perkins, "IP Encapsulation Within IP," IETF Network Working Group, RFC 2003 (October 1996), which is incorporated by reference herein. The original IP packet includes a header and a data region which constitutes the payload of the IP packet. To encapsulate an IP packet using IP in IP encapsulation, an outer IP header 410 is inserted before the packet's existing header 420 and data region 405. The outer IP header Source IP Address 403 and Destination IP Address 404 identify the endpoints of the tunnel. The original IP packet remains basically unchanged and can be readily de-capsulated by stripping off the outer IP header 410. FIG. 5 is a conceptual representation of another process of encapsulating a standard IPv4 packet, in accordance with the technique known in the art as "Minimal Encapsulation Within IP." See C. Perkins, "Minimal Encapsulation Within IP," IETF Network Working Group, RFC 2004 (October 1996), which is incorporated by reference herein. Minimal encapsulation achieves the same objective as above without actually including the full IP header of the original packet. Instead, the Destination IP Address field 502 of the original packet (and optionally the Source IP Address field 501) is modified to force routing to the intermediate destination—Destination IP Address 504—and the protocol field is modified to indicate that minimal encapsulation is being utilized. The original values of the Destination IP Address field (and the Source IP Address field if needed) and protocol field are saved in an eight or twelve octet extension to the original IP header. The extension and the header are depicted as 520 and 510, respectively, in FIG. 5. The Total Length field of the original packet is also changed to reflect the expanded length, and the Header Checksum recalculated. When the intermediate destination receives the packet, it observes the "minimal encapsulation" protocol field, restores the original packet based on the values carried in the extension 520 to the IP header 510, and forwards the

original packet. Where minimal encapsulation is initiated by an intermediate router, perhaps based on a packet filter, both Source and Destination IP Address fields are modified in the original packet and carried in a twelve octet header extension. If minimal encapsulation is initiated by the source of the packet, there is no need to modify the Source IP Address field in the IP header, and an eight octet minimal encapsulation extension can be added to the IP header. The format of the minimal encapsulation extension 520 to the IP header 510 is shown in FIG. 5. The fields in the extension are defined in RFC 2004 as follows:

PROTOCOL	Copied from the Protocol field in the original IP header.
ORIGINAL SOURCE ADDRESS PRESENT (S) (508 in FIG. 5)	<p>0 The Original Source Address field is not present. The length of the minimal tunneling header in this case is 8 octets.</p> <p>1 The Original Source Address field is present. The length of the minimal tunneling header in this case is 12 octets.</p>
RESERVED	Sent as zero; ignored on reception.
HEADER CHECKSUM	The 16-bit one's complement of the one's complement sum of all 16-bit words in the minimal forwarding header. For purposes of computing the checksum, the value of the checksum field is 0. The IP header and IP payload (after the minimal forwarding header) are not included in this checksum computation.
ORIGINAL DESTINATION ADDRESS (502 in FIG. 5)	Copied from the Destination Address field in the original IP header.
ORIGINAL SOURCE ADDRESS (501 in FIG. 5)	Copied from the Source Address field in the original IP header. This field is present only if the Original Source Address Present (S) bit is set.

195 FIG. 6 and FIG. 7 set forth the processing performed at a network access device and a service network router, respectively, to force service-related packets to route through the service network. The network access device has been pre-configured with its IP address, with its service-related IP address, and with the IP address of the service network router at the other end of the tunnel. It is

200 advantageous to provide a service activation system which permits the dynamic allocation, assignment, and reassignment of the IP addresses to the plurality of network access devices based on customer subscriptions to particular services, as further described in copending utility patent application, "SERVICE SELECTION IN A SHARED ACCESS NETWORK USING DYNAMIC HOST

205 CONFIGURATION PROTOCOL," filed contemporaneously with the present application, and incorporated by reference herein. At step 601, the process running on the network access device (or alternatively on another device on behalf of the network access device) receives a packet that has been constructed by another process or another part of the same process. At step 602, the packet is

210 determined to be service-related and, thus, outbound to the relevant service network. The packet has the service-related IP address in the source address field and a destination address. At step 603, the packet is encapsulated using, for example, the encapsulation techniques described above. The destination address field of the new packet is the IP address of the service network router. At step

215 604, the encapsulated packet is tunneled to the service network router in the service network. FIG. 7 sets forth the processing performed at the router in the service network. At step 701, the router receives an incoming packet. At step 702, the router reads the packet header and retrieves any packet filtering rules reflected in access lists and the rules provided for encapsulation and de-

220 encapsulation of packets. At step 703, the router compares the destination address to its own address (or the address of a virtual interface, as further described herein) and determines whether the packet has been encapsulated. At step 704, the router decapsulates the packet and routes the packet to the original destination address field in the packet. Before decapsulating the packet, the router may check

225 through a list of authorized service-related IP addresses to ensure that the packet comes from a properly authenticated subscriber's network access device. The use of encapsulation by the network access device and the service network router is transparent to the access network infrastructure.

230 Packets directed to the subscriber from the rest of the Internet can be addressed to the tunneling IP address of the network access device and forwarded using normal routing procedures without tunneling. Where for some reason the service provider wishes to route service-related traffic back through the service network, such packets can be addressed to the service-related IP address and routed back to the service network router. With reference to FIG. 7 again, the service network router receives the incoming packet at step 701 and determines at 235 step 705 that the destination address of the packet matches a service-related IP address associated with a particular subscriber. At step 706, the router encapsulates the packet using, for example, an encapsulation technique describe above. The router accesses a database of service subscribers and determines the 240 tunneling IP address associated with the service-related IP address used by the particular subscriber. The router then uses this IP address in the destination field when encapsulating the packet. At step 707, the packet is tunneled to the network access device associated with the subscriber. With reference to FIG. 6 again, the network access device process receives the packet at step 601 and determines at 245 step 605 that the packet is encapsulated and from the encapsulating router in the service network. At step 606, the packet is decapsulated and processed accordingly by the relevant application.

Any communication between the network access device and devices attached to the access network infrastructure need not use encapsulation. 250 Such packets can be processed normally by the network access device at step 607 using the non-service-related IP address and routed by the access network infrastructure using normal routing procedures. Note that communications between the network access device and the service network provider itself, e.g. a DNS query, do not need to use tunneling – but could use tunneling depending on 255 the needs of the different entities.

The service provider router that de-encapsulates packets may be a single point of failure that may block customer access to service provider services. It is advantageous to provide procedures to eliminate this single point of failure. The address used by subscribers to forward packets should be routed to an available router in the service network. One method of accomplishing this is to have the operator of the service network choose a subnet to provide the address of the de-encapsulation router(s). Each de-encapsulation router is configured with a virtual interface with an address on the specified subnet. See, e.g., C. Perkins, "IP Mobility Support," IETF Network Working Group, RFC 2002 (October 1996), which is incorporated by reference herein. No "real" interfaces are allowed on this subnet. Note that a "virtual interface" exists at an IP level for routing purposes, but is not associated with any physical port. Each router advertises connectivity to the specified subnet, but since all interfaces on the subnet are virtual interfaces, there is no local connectivity (via the specified subnet). Routers on this "virtual subnet" need not be "close" to each other in a routing sense. The directed subnet broadcast address (the host part of the address can be all ones) looks like a normal host address to every router that does not know the subnet mask. Therefore a packet addressed to a directed subnet broadcast address will be forwarded to the "closest" router advertising connectivity to the subnet. If a subscriber's network access device uses the directed subnet broadcast address of the subnet shared by de-encapsulation routers as the destination of encapsulated packets, normal routing procedures will forward the packet to the "closest" router advertising connectivity to the subnet. The de-encapsulation router that receives the packet will recognize that the packet is a broadcast, and process the packet (since it has an address on the subnet). Since the interface associated with the subnet is a virtual interface, the router cannot forward the packet to other members of the subnet. Normal routing procedures will ensure that packets are forwarded to the "closest" available de-encapsulation router, advantageously making appropriate adjustments as routers fail and/or recover from failure.

The above embodiments have been described from the perspective of layer three tunneling. Another embodiment of this aspect of the invention is to

use a layer two tunneling technique between the network access device and a service network node acting as a layer two tunnel termination device. For example, FIG. 3A sets forth another exemplary access architecture based on an HFC network, roughly corresponding to FIG. 2, using the Layer Two Tunneling Protocol (L2TP). See, e.g., W. Townsley, A. Valencia, G. Zorn, A. Rubens, G. Pall, B. Palter, "Layer Two Tunneling Protocol (L2TP)," IETF Network Working Group, IETF draft, draft-ietf-l2tp-l2tpbis-01.txt (November 2000). Rather than using layer three routers in FIG. 2, the service networks 351 and 352 in FIG. 3 have L2TP Network Servers ("LNS") 341 and 342. An LNS can be a router or other network node configured to act as a layer two tunnel terminating device. FIG. 3B is a conceptual diagram of the end-to-end communication protocol stack from a network access device 301 to an LNS 341 in the service provider's network 351 where L2TP is utilized. The network access device encapsulates PPP packets in L2TP for transport across the access network infrastructure to the relevant service network. The PPP packets are tunneled across the access network infrastructure to an LNS in the service network which strips the L2TP and terminates PPP.

The foregoing Detailed Description is to be understood as being in every respect illustrative and exemplary, but not restrictive, and the scope of the invention disclosed herein is not to be determined from the Detailed Description, but rather from the claims as interpreted according to the full breadth permitted by the patent laws. It is to be understood that the embodiments shown and described herein are only illustrative of the principles of the present invention and that various modifications may be implemented by those skilled in the art without departing from the scope and spirit of the invention. For example, the detailed description describes an embodiment of the invention with particular reference to an HFC access network architecture. However, the principles of the present invention could be readily extended to other access network architectures, such as DSL, wireless, satellite, etc. Such an extension could be readily implemented by one of ordinary skill in the art given the above disclosure.